

How High a Degree is High Enough for High Order Finite Elements?

William F. Mitchell *

National Institute of Standards and Technology, Gaithersburg, Maryland, U.S.A.
william.mitchell@nist.gov

Abstract

High order finite element methods can solve partial differential equations more efficiently than low order methods. But how large of a polynomial degree is beneficial? This paper addresses that question through a case study of three problems representing problems with smooth solutions, problems with steep gradients, and problems with singularities. It also contrasts h -adaptive, p -adaptive, and hp -adaptive refinement. The results indicate that for low accuracy requirements, like 1 % relative error, h -adaptive refinement with relatively low order elements is sufficient, and for high accuracy requirements, p -adaptive refinement is best for smooth problems and hp -adaptive refinement with elements up to about 10^{th} degree is best for other problems.

Keywords: finite elements, high order methods, hp-FEM

1 Introduction

With fast convergence rates, high order finite element methods (FEM) can solve partial differential equations more efficiently in terms of accuracy vs. computational resources than low order methods. For problems in which the solution exhibits some local phenomenon, like a wave front or a singularity, use of adaptive mesh refinement further improves the efficiency. There are three versions of these finite element methods: 1) h -adaptive finite elements, in which accuracy is improved by decreasing the size, h , of selected elements while keeping a fixed polynomial degree p , 2) p -adaptive finite elements, in which accuracy is improved by increasing the polynomial degree of selected elements while keeping the same spacial grid, and 3) hp -adaptive finite elements, in which both element size and polynomial degree are changed.

But how large of a polynomial degree is beneficial? As illustrated in Figure 2 of Section 4, increasing the degree by one when it is small has a dramatic effect on the convergence rate, but once the degree becomes fairly large, the convergence rate does not change much. Moreover, high order elements are more expensive computationally because 1) higher degree polynomial

*Preprint. Appeared in *Procedia Computer Science*, 15, 2015, pp. 246–255.

basis functions require more operations to compute, 2) the quadrature rules for numerical integration to compute the matrix for the discrete problem must be of correspondingly higher order, which means more quadrature points at which to evaluate the basis functions, and 3) the matrix for the discrete problem is denser, which means more computation to solve the linear system. One would suspect that at some point the trade-off of more computation for a small gain in convergence rate is no longer beneficial.

This paper examines this question through a case study involving three 2D test problems of different character – a smooth solution, a steep gradient, and a singularity. In each case, the three versions of the finite element method are used with different polynomial degrees to determine when increasing the degree no longer pays off.

The remainder of the paper is organized as follows. Section 2 describes the three test problems used in the case study. Section 3 gives the details of the finite element methods used. The main section of the paper is Section 4 which presents and discusses the results of the computations. Finally, the conclusions are presented in Section 5.

2 Test Problems

Three test problems from the NIST Adaptive Mesh Refinement Benchmark Problems [?] are used to study the convergence properties of the finite element methods.

Mild wave front. The arctan circular wave front problem is Poisson’s equation on the unit square with Dirichlet boundary conditions:

$$\begin{aligned} -\frac{\partial^2 u}{\partial x^2} - \frac{\partial^2 u}{\partial y^2} &= f \text{ in } \Omega = (0, 1) \times (0, 1) \\ u &= g \text{ on } \partial\Omega \end{aligned}$$

The functions f and g are chosen so that the exact solution is $\tan^{-1}(\alpha(r - r_0))$ where $r = \sqrt{(x - x_0)^2 + (y - y_0)^2}$. The mild wave front uses $(x_0, y_0) = (-0.05, -0.05)$, $r_0 = 0.7$ and $\alpha = 20$. α controls the steepness of the wave, i.e., the magnitude of the gradient within the wave front. The solution for this problem is very smooth, so conventional wisdom is that p -refinement and large values of p should be appropriate.

Steep wave front. This is the same as the mild wave front problem, but with $\alpha = 1000$. This is a very steep wave front with a rapid transition from a gradient near zero to a very large gradient.

L-shaped domain. This is Laplace’s equation (Poisson’s equation with $f = 0$) on an L-shaped domain given by $(-1, 1) \times (-1, 1) \setminus [0, 1) \times (-1, 0]$, i.e. a square with the lower right quadrant removed. The exact solution is $r^{2/3} \sin(2\theta/3)$ where $r = \sqrt{x^2 + y^2}$ and $\theta = \tan^{-1}(y/x)$. This has a point singularity at the origin. Conventional wisdom is to perform h -refinement at the singularity and p -refinement elsewhere.

3 Details of the Finite Element Methods

The usual continuous Galerkin finite element method is used with triangular elements. h -refinement is by newest node bisection [?]. The method begins with the coarse grid shown in




Figure 1: Initial grid for the wave front problems.

Figure 1 (the grid for the L-shaped domain problem is similar) and alternates between phases of refinement and solution of the discrete problem. A refinement phase consists of refining all elements for which the error estimator is larger than half of the largest error estimator. The error estimator is a high order version of the Bank-Weiser error estimator [?]. The basis functions for the finite element space are the Szabo-Babuška p -hierarchical basis functions [?]. A multigrid solver is used to solve the linear system of the discrete problem.

All three versions of the finite element method are used. For h -adaptive FEM, a uniform degree, p , is used for each p between 1 and 21. For p -adaptive FEM, a maximum degree of 21 is imposed. If refinement is indicated for an element of degree 21, that element is not refined. If all elements that are marked for refinement have degree 21, refinement stalls and the program terminates. For hp -adaptive FEM, a maximum degree, p_{\max} , is specified for each value between 1 and 21. When refinement is indicated for an element of degree p_{\max} , it is refined by h regardless of the type of refinement indicated by the hp -adaptive strategy.

For the two wave front problems, the “type parameter” hp -adaptive strategy [?] is used. This strategy performed well on the circular wave front problems in a recent study of hp -adaptive strategies [?]. For the L-shaped domain problem, the *a priori* knowledge of the singularity is used for the hp -adaptive strategy. Elements that touch the origin are refined by h , and all others are refined by p .

4 Numerical Results

Computations were performed on one core of an Intel Xeon 2.5 GHz EMT64 ¹ operating under Red Hat Enterprise Linux release 5.10. The finite element code PHAML version 1.14.0 [?] was compiled with the Intel Fortran compiler version 15.0.0.

The convergence of each method is examined via the relative energy norm of the error

$$\left(\frac{\int_{\Omega} \left(\frac{\partial(u-\tilde{u})}{\partial x} \right)^2 + \left(\frac{\partial(u-\tilde{u})}{\partial y} \right)^2}{\int_{\Omega} \left(\frac{\partial u}{\partial x} \right)^2 + \left(\frac{\partial u}{\partial y} \right)^2} \right)^{1/2}$$

¹The mention of specific products, trademarks, or brand names is for purposes of identification only. Such mention is not to be interpreted in any way as an endorsement or certification of such products or brands by the National Institute of Standards and Technology. All trademarks mentioned herein belong to their respective owners.




Figure 2: Convergence graph of relative energy norm of the error vs. degrees of freedom for hp -adaptive refinement with the mild wave front problem.

vs. the number of degrees of freedom, and vs. the computation time, where \tilde{u} is the finite element approximation to u . Two accuracy requirements are considered. The low accuracy of 10^{-2} (i.e. 1 % error) is typical of engineering applications. The high accuracy of 10^{-6} arises in some scientific applications.

Figure 2 illustrates a typical convergence graph. This is the graph for hp -adaptive refinement with the mild wave front problem. One can see that the convergence curves become nearly identical once p_{\max} is sufficiently large, about $p_{\max} = 10$ in this example. One can also see that, for a given p_{\max} , convergence of the hp -adaptive FEM is initially exponential, as indicated by the curvature of the convergence curve, but eventually transitions to algebraic, as indicated by the straightness of the convergence curve. The algebraic convergence corresponds to the convergence rate of h -adaptive FEM with $p = p_{\max}$, and indicates that the refinement is eventually dominated by h -refinement of elements of degree p_{\max} .

The tables of Sections 4.1-4.3 (see, for example, Table 2, which corresponds to Figure 3) give the largest value of p (or p_{\max} , throughout) that is beneficial under several criteria. They give the result for low and high accuracy, and for error vs. degrees of freedom (dof) and error vs. computation time. The value of p is determined as illustrated in Figure 3. The value listed as optimal is the smallest value of p that gives minimum dof (or time, throughout), 13 in the illustration. One may also consider that the optimal p is not necessary, just one that is close enough that it produces not much more work. The column labeled 10 % is the smallest value of p for which the dof is less than 10 % more than the dof of the optimal p , 10 in the illustration. Similarly the 2X column gives the p for which the dof is less than twice the optimum, 6 in the illustration.




Figure 3: Determining the optimal, 10 %, and 2X values of p_{\max} for the high accuracy solution of the mild wave front problem.

| accuracy | degrees of freedom | | | time | | |
|-----------|--------------------|------|----|---------|------|----|
| | optimal | 10 % | 2X | optimal | 10 % | 2X |
| 10^{-2} | 3 | 3 | 2 | 4 | 4 | 4 |
| 10^{-6} | 8 | 7 | 6 | 12 | 12 | 6 |

Table 1: Largest beneficial value of p for the mild wave front problem with h -adaptive refinement.

4.1 Mild Wave Front

Table 1 gives the largest beneficial p for h -adaptive refinement for the mild wave front problem. For low accuracy, degrees 3 and 4 are optimal for dof and time, respectively. For high accuracy solutions, higher order elements are more effective, with degree 8 for dof and degree 12 for time. In all cases, there is little, if any, reduction in p by allowing a 10 % increase in dof or time. Allowing twice the resources also has little, if any, reduction in p , except for error vs. time for the high accuracy solution where there is a substantial decrease from 12 to 6.

It may seem surprising that, for the high accuracy solution, a much higher degree is optimal for error vs. time than that for error vs. dof, given the increase in the amount of computation per degree of freedom as the polynomial degree is increased. The explanation is that the degree 12 run used many fewer refine-solve loops than the degree 8 run (21 vs. 40), which means many fewer instances of assembling and solving the linear system of the discrete problem. The reason for fewer loops is that the higher order approximation has a smaller error on the initial grid, 1.24×10^{-5} as opposed to 2.89×10^{-4} .

Table 2 contains the largest beneficial p_{\max} for hp -adaptive refinement with the mild wave problem. Again degree 3 or 4 is optimal for low accuracy solutions. And again allowing higher degrees is beneficial for the high accuracy solution. But in this case the higher degree does not

| accuracy | degrees of freedom | | | time | | |
|-----------|--------------------|------|----|---------|------|----|
| | optimal | 10 % | 2X | optimal | 10 % | 2X |
| 10^{-2} | 4 | 3 | 2 | 3 | 3 | 2 |
| 10^{-6} | 13 | 10 | 6 | 7 | 6 | 5 |

Table 2: Largest beneficial value of p_{\max} for the mild wave front problem with hp -adaptive refinement.

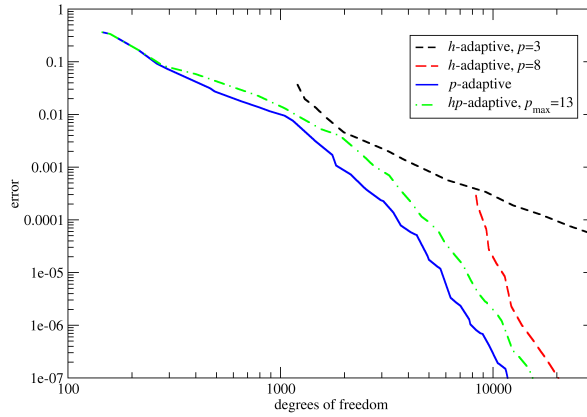


Figure 4: Error vs. degrees of freedom convergence curves for h -adaptive refinement, p -adaptive refinement, and hp -adaptive refinement for the mild wave front problem.

benefit from a more accurate solution on the initial grid (all of them start with $p = 1$), so the optimal p_{\max} is considerably smaller for error vs. time than for error vs. dof. But, the rate of improvement in error vs. dof by increasing p_{\max} drops rapidly after $p_{\max} = 6$ so that if one allows doubling the resources there is only a difference of one between the largest beneficial p_{\max} for dof and time.

Figures 4 and 5 compare the three types of finite element methods, h -adaptive, p -adaptive and hp -adaptive, using the optimal p and p_{\max} from Tables 1 and 2. hp -adaptive with $p_{\max} = 8$ is the same as $p_{\max} = 12$ in the vicinity of error = 10^{-2} , so that curve is not included. For error vs. dof, hp -adaptive refinement is superior to h -refinement with the optimal p , but p -adaptive refinement is even better. For error vs. time, p -adaptive refinement is still superior to hp -adaptive refinement, but h -adaptive refinement is better than p -adaptive refinement, especially at low accuracy. This is probably because the p -adaptive refinement has higher order elements than the selected fixed degree for h -adaptive refinement. For this example they were as high as $p = 6$ for error = 10^{-2} and $p = 17$ for error = 10^{-6} . But for h -adaptivity to be superior to p -adaptivity, one must know the optimal p for the desired accuracy.

4.2 Steep Wave Front

Table 3 gives the largest beneficial p for h -adaptive refinement for the steep wave front problem. At low accuracy, $p = 4$ is optimal for error vs. dof, and $p = 6$ is optimal for error vs. time.




Figure 5: Error vs. computation time convergence curves for h -adaptive refinement, p -adaptive refinement, and hp -adaptive refinement for the mild wave front problem.

| accuracy | degrees of freedom | | | time | | |
|-----------|--------------------|------|----|---------|------|----|
| | optimal | 10 % | 2X | optimal | 10 % | 2X |
| 10^{-2} | 4 | 3 | 3 | 6 | 4 | 3 |
| 10^{-6} | 11 | 9 | 6 | 8 | 6 | 5 |

Table 3: Largest beneficial value of p for the steep wave front problem with h -adaptive refinement.

However, with just a 10 % increase in time, $p = 4$ is sufficient. At high accuracy, it is beneficial to use higher order elements.

Table 4 gives the largest beneficial p_{\max} for hp -adaptive refinement for the steep wave front problem. For the low accuracy solution, $p_{\max} = 9$ is optimal for error vs. dof, but with a 10 % increase in the dof $p_{\max} = 4$ is sufficient. The high accuracy solution benefits from very high order elements, with $p_{\max} = 18$ optimal for error vs. dof, but allowing twice as many dof reduces it to $p_{\max} = 7$, the same as the optimal p_{\max} for error vs. time.

Figures 6 and 7 compare the three types of finite element methods. p -adaptive refinement stalls and is ineffective on this problem, so it is omitted from Figure 7. The steep wave front cannot be resolved by high degree polynomials alone; small elements are required. hp -adaptive refinement is better than h -adaptive refinement in terms of error vs. dof, but h -adaptive refinement is better in terms of time if the correct value of p is used.

4.3 L-shaped Domain

Table 5 gives the largest beneficial p for h -adaptive refinement for the L-shaped domain problem. Low order $p = 2$ elements are most effective for low accuracy in error vs. dof, but in error vs. time $p = 5$ is optimal, with $p = 2$ within a factor of two. For high order, $p = 5$ or 6 is most effective.

Table 6 gives the largest beneficial p_{\max} for hp -adaptive refinement for the L-shaped domain problem. For low accuracy with an allowance of 10 % above the optimal, $p = 2$ is sufficient.

| accuracy | degrees of freedom | | | time | | |
|-----------|--------------------|------|----|---------|------|----|
| | optimal | 10 % | 2X | optimal | 10 % | 2X |
| 10^{-2} | 9 | 4 | 3 | 4 | 4 | 2 |
| 10^{-6} | 18 | 14 | 7 | 7 | 6 | 5 |

Table 4: Largest beneficial value of p_{\max} for the steep wave front problem with hp -adaptive refinement.




Figure 6: Error vs. degrees of freedom convergence curves for h -adaptive refinement, p -adaptive refinement, and hp -adaptive refinement for the steep wave front problem.

High accuracy benefits from higher order elements with $p_{\max} = 8$ or 9 .

Figures 8 and 9 compare the three types of finite element methods. p -adaptive refinement stalls and is not effective, because high order elements are not good at approximating singularities. At low accuracy there is effectively no difference between h -adaptive refinement and hp -adaptive refinement in error vs. dof, but h -adaptive refinement is considerably faster. At high accuracy, hp -adaptive refinement is superior in both dof and time.

5 Conclusion

This paper addressed the question of how large of a polynomial degree is beneficial to high order finite element methods, and contrasted h -adaptive, p -adaptive and hp -adaptive refinement. Three test problems were used representing problems with smooth solutions, problems with steep gradients, and problems with singularities.

In general it appears that for a low accuracy requirement like 1 % relative error, h -adaptive refinement with relatively low order elements is sufficient. The singular problem did not benefit from $p > 2$, and for the nonsingular problems $p = 3$ or 4 is large enough.

For high accuracy requirement, higher order elements are appropriate. But there is little benefit in going past $p = 10$ or so in the convergence with respect to the number of degrees of freedom, and a maximum p of about 6 or 8 is sufficient in terms of optimal computation time,




Figure 7: Error vs. computation time convergence curves for h -adaptive refinement and hp -adaptive refinement for the steep wave front problem.

| accuracy | degrees of freedom | | | time | | |
|-----------|--------------------|------|----|---------|------|----|
| | optimal | 10 % | 2X | optimal | 10 % | 2X |
| 10^{-2} | 2 | 2 | 2 | 5 | 4 | 2 |
| 10^{-6} | 6 | 6 | 5 | 5 | 5 | 4 |

Table 5: Largest beneficial value of p for the L-domain problem with h -adaptive refinement.

except for smooth problems. For smooth problems, p -adaptive refinement is most effective and is recommended. For other problems, hp -adaptive refinement is superior to h -adaptive refinement in terms of error vs. dof. h -adaptive refinement was faster, but only if the optimal p for the given accuracy requirement is known, so hp -adaptive refinement is recommended.

| accuracy | degrees of freedom | | | time | | |
|-----------|--------------------|------|----|---------|------|----|
| | optimal | 10 % | 2X | optimal | 10 % | 2X |
| 10^{-2} | 3 | 2 | 2 | 8 | 2 | 2 |
| 10^{-6} | 9 | 8 | 5 | 8 | 6 | 4 |

Table 6: Largest beneficial value of p_{\max} for the L-domain problem with hp -adaptive refinement.




Figure 8: Error vs. degrees of freedom convergence curves for h -adaptive refinement, p -adaptive refinement, and hp -adaptive refinement for the L-shaped domain problem.




Figure 9: Error vs. computation time convergence curves for h -adaptive refinement and hp -adaptive refinement for the L-shaped domain problem.